THE EUROPEAN MOONSHOT FOR EDUCATION

AI-Powered Learning as Sovereign Infrastructure

A proposal for the twelfth Moonshot project
in the EU Multiannual Financial Framework 2028–2034

Matthias Pochmann

2026

DOI: 10.5281/zenodo.18759299

Companion paper: A Generative Education Architecture
for Planetary-Scale Personalized Learning (Pochmann 2026)

AI in education is an infrastructure question, and the infrastructure has already been decided — just not by any government. In mid-2025, Google embedded LearnLM into the productivity suite already used by 170 million students and educators across more than 230 countries and territories. No learner opted in. No pedagogical authority reviewed it. A school administrator toggled a setting, and the largest AI tutoring deployment in history was live.

The MFF 2028–2034 proposal lists eleven Moonshot projects — fusion energy, quantum computing, clean aviation, next-generation AI, data sovereignty, automated transport, regenerative therapies, space economy among them. The system that trains the scientists, engineers, and physicians for all eleven — the floor on which every other Moonshot stands — is absent. Education is not on the list.

A companion paper describes the architecture for a European alternative in full (Pochmann 2026). The question this document answers is not how such a system works, but who decides whether it gets built — and on what terms.

Every Moonshot on the list needs trained minds. The system that produces those minds is not on the list. Education as a twelfth Moonshot project in the MFF 2028–2034 — not as an Erasmus extension, but as infrastructure on the level of CERN, fusion energy, and quantum computing.

The Fork in the Road

Google did not need to win 170 million users. It had them. Google Workspace for Education is the default operating layer in school systems across more than 230 countries and territories — assignments, grading, communication, identity, all running through a single platform. The decision to adopt that platform was institutional and often deliberate — involving procurement processes, data processing agreements, and in Europe, DPIA assessments. But in mid-2025, LearnLM arrived as a feature update within that platform. No new procurement decision. No new pedagogical review. No new DPIA. An IT administrator toggled a setting in a console, and AI tutoring was live — on every account the district controlled. The AI activation decision was architecturally embedded in a prior platform decision, bypassing the governance mechanisms that would normally apply to a new educational technology deployment.

> **This is cognitive lock-in — not adoption, but activation.**

The engineering is serious. In controlled studies, students using LearnLM under human tutor supervision were 5.5 percentage points more likely to solve novel problems than those working with human tutors alone. The results are real, and the deployment at scale is an accomplishment. But the architecture

decision that embedded this technology into a generation's daily learning was made without anyone deciding.

The OECD has measured what happens when general-purpose AI enters education without intentional pedagogical architecture. The Digital Education Outlook 2026 found that students with access to general-purpose generative AI produce higher-quality outputs — better-structured essays, more complete problem sets. But when those students sit exams without AI access, the advantage disappears. In some cases it reverses. The OECD's language is precise: "Offloading cognitive tasks to general-purpose chatbots creates risks of metacognitive laziness and disengagement."

> **The students performed more. They did not learn more.**

Purpose-built educational AI with intentional pedagogical design — tools that scaffold thinking rather than replace it — shows sustained improvements that persist when the AI is removed. LearnLM is closer to this second category than to a general-purpose chatbot; its Socratic scaffolding is real, and its measured learning gains are genuine. But the OECD finding establishes the baseline risk of the ecosystem into which 170 million accounts have been activated: most AI that students encounter remains general-purpose, and even purpose-built tools operate within structural constraints that limit not pedagogical quality but institutional form. What is activated on 170 million accounts is architecturally incomplete — not because the tutoring is ineffective, but because effective tutoring without data sovereignty, longitudinal profiles, auditable knowledge governance, institutional permanence, and offline capability is a partial answer to a structural question.

If digital sovereignty is a European value — and it is; the Commission's 2026 Work Programme carries the title "Europe's Independence Moment" — then the chain of reasoning is short. GDPR was the answer to data misuse. The AI Act was the answer to uncontrolled artificial intelligence. Moonshot projects are the answer to strategic infrastructure gaps. Chip sovereignty, data sovereignty, digital sovereignty — each recognized a domain where dependence on external architectures posed a structural risk. Cognitive sovereignty is the next link in that chain: control over the systems that shape how a generation learns to think, what knowledge they encounter, and who holds the record of their intellectual development. The question is not whether Europe needs a sovereign education infrastructure. It is why education is absent from the Moonshot list.

> **Your children already learn with an AI that does not answer to you, does not explain itself to you, and can change its terms tomorrow.**

Google builds good technology. The problem is not capability. It is institutional form. Five structural constraints define what no platform company — regardless of engineering talent — can deliver:

A policy promise that student data will not train models is reversible by a terms-of-service update; an architecture where the cognitive profile is processed only on the learner's local system, with no interface through which it could be bulk-extracted or queried by third parties, requires a fiduciary structure no company whose revenue depends on data aggregation can credibly adopt. The real distinction is not whether the system processes the data — any system that personalizes must process it — but whether any entity beyond the learner's trust boundary can extract, aggregate, or repurpose it. Google tracks what a student clicked and how long they hesitated; it does not build a cognitive model that accompanies a person from childhood literacy through adult reskilling — and when the student leaves, they take nothing. Neither Google nor Khan Academy maintains a structured, auditable knowledge graph separating universal scientific consensus from culturally sovereign content; knowledge definition remains proprietary. A publicly traded corporation has fiduciary duties to shareholders; if LearnLM becomes premium, is discontinued, or changes its pedagogical logic, no governance body prevents it. Everything is cloud-based, requiring accounts and stable internet — for 272 million children without school, many without reliable electricity, these tools do not exist.

Whoever curates the knowledge graph curates the future.

These are not failures of execution. They are consequences of structure. A company that monetizes attention cannot build a system designed to make itself unnecessary. A company governed by shareholder return cannot place a learner's cognitive profile beyond its own reach. A company that operates in the cloud cannot serve a child who has no connection. The gap between what platform products deliver and what trustee infrastructure requires is not a quality gap. It is a category difference.

Four forces make the timing irreversible. Inference costs for GPT-3.5-level performance fell 280-fold in under two years (Stanford HAI AI Index 2025). Edge AI hardware performance doubles every 1.9 years on the same form factor; quantized models under nine billion parameters already run on smartphones, and if the trend holds, phones will run ten-billion-parameter models natively before the end of the decade. Satellite connectivity from competing constellations reaches global coverage by 2028 at $5–10 per month in developing countries. And 272 million children remain out of school — a figure largely unchanged since 2015 — while 300 million currently in school

will leave without basic literacy. Each curve is independently documented. Laid side by side, they trace a single trajectory: AI-powered education for every learner on Earth will arrive — as inevitably as open standards replaced proprietary fragmentation when GSM unified a continent's networks.

MOOCs promised democratization; completion rates settled at 5–15 percent. Intelligent Tutoring Systems remained niche — each covering a single narrow subject, requiring hundreds of development hours per hour of content. The difference is structural: MOOCs digitized distribution. ITS digitized tutoring. Generative AI changes what is *produced* — the learning experience itself, generated for this person at this moment. That is not an incremental improvement. It is a different category of system.

Every year the decision is deferred, the installed base grows by tens of millions. Twelve-year-olds entering the system today will complete their entire educational biography inside a single company's ecosystem before anyone builds an alternative. The technology that makes AI-powered education universal is arriving regardless of any policy decision. Whether that education runs on trustee infrastructure — where the learner owns the profile, the knowledge graph is auditable, and governance answers to no shareholder — or on platform products where cognitive development is a feature that can be repriced, restructured, or revoked: that is the fork. Architecture decisions, once the installed base is built, do not reverse.

> **The generation now in school will either own its learning or rent it. That choice is being made — by default, in admin consoles, one toggle at a time.**

The Architecture That Emerges

She is twelve, lives in rural Rwanda, and her bridge collapsed. Not a real bridge — a simulated one, built from beams she sized and joints she placed, loaded until something gave. She says: "I want to understand why." The Mentor does not explain. It rebuilds the bridge and hands her the lever arm. She shortens it. Forces shift. Stress concentrations bloom red at the joint she weakened. The simulation is not a video about physics. It is physics. She is not watching. She is inside.

The environment was generated for this person, at this moment, in response to this failure. No curriculum author wrote it. No content library contained it. When she asked her question, no material existed that answered it — for her, in her state of understanding, with her specific misconception about where load concentrates. The Mentor created it: a simulation with physical parameters

she can change, forces she can see, consequences she can feel before she can name them. She does not know she is doing structural engineering. She knows her bridge broke, and she wants to know why. That is enough. That has always been enough.

In Thessaloniki, a forty-year-old man sits in a repurposed shipping container fitted with local compute. His factory closed eight months ago. The system knows he thinks spatially, learned fractions through cooking ratios, stalls on abstract notation, responds to physical metaphor — a profile built from thousands of micro-cognitive signals accumulated over months of daily interaction. Today the Mentor opens with a robotic arm on screen, joints highlighted. "Move this one," it says. He drags the shoulder joint. The arm sweeps wide. "Now lock the elbow and move the shoulder again." The sweep tightens. He sees it before the Mentor names it: constrained degrees of freedom change what a joint can reach. The kinematic chain is not an equation on a whiteboard. It is a thing he moved with his hands and watched behave — because the system knew that for him, understanding starts in the body and arrives at notation later. Same destination as every robotics student. A different route — not because the standards differ, but because the path to the same standard runs through who he already is. No standardized curriculum has ever addressed this: the path to the same competence is as individual as the mind walking it.

The profile that guides this route is not a preference setting. It is a compression of how his mind works: where he hesitates, what metaphors unlock understanding, which abstractions resist his first approach and yield to his second. No employer sees it. No government agency can request it. Not because a regulation forbids access — because the architecture ensures that no entity other than the learner's local system processes the cognitive profile, and no interface exists through which it could be bulk-extracted or queried by third parties. The most intimate data about a human being is not a medical history. It is the record of how their mind learns — where it fails, what it fears, how it restructures itself after confusion. That record exists today, scattered across platforms whose terms of service grant themselves perpetual license to it. Here, it belongs to him. When he leaves, it goes with him.

In São Paulo, a fifteen-year-old opens her session. The Mentor speaks first: "November 12th — you cracked vector calculus. The same pattern shows up today." She remembers. The Mentor remembers too — not as a data point retrieved from a log, but as a curated moment it recognized as significant and preserved. It draws a line between that breakthrough and the problem in front of her now, a structural resonance she would not have seen on her own.

But she is not alone with a screen. She is in a hub — a physical learning space in her neighborhood, converted from a community center, with workstations

and a shared simulation wall. The system has matched her with two other learners: one in Nairobi, working on a complementary aspect of the same engineering problem, and one sitting three meters away. The AI translates in real time — not just Portuguese to Swahili but context: when the student in Nairobi references a construction technique common in East Africa, the system bridges the gap so the metaphor lands. They are not consuming the same content in parallel. They are solving different faces of the same problem, and the solution requires all three perspectives. The hub is loud. Someone is laughing at a failed simulation on the other side of the room. A younger child is reading in a corner, headphones on, in a session entirely her own. The system that individualized the twelve-year-old's bridge and the forty-year-old's retraining path now does something no optimized feed has ever done: it finds the person whose thinking complements yours, and it puts you to work together. Platform architecture individualizes into isolation — each user alone before a stream refined to maximize engagement. This architecture individualizes into connection. The hub is the physical proof that infrastructure creates places. Platforms create screens.

Three scenes. Three architectures made visible. What connects them is not a feature set but a structural claim: education is not content delivery. It is the construction of a mind — and the construction of a mind requires that the learner, not the platform, holds the blueprint.

The components exist at varying levels of maturity, each proven in isolation. Adaptive platforms demonstrated decades ago that dynamic adjustment to learner performance works. Privacy-preserving computation has been demonstrated commercially through multiple approaches, each with known limitations. Edge AI and federated learning reduce the attack surface by keeping raw data local, but no existing integration achieves the full guarantee at scale. Data cooperatives — most notably in healthcare — have proven that citizens can govern sensitive personal data through democratic structures while preserving the data's utility for the collective. Edge AI crossed a critical threshold: quantized models run on smartphones, and on-device fine-tuning enables learner-specific behavior without exporting a single data point. Federated knowledge graph completion has shown that joint learning across multiple graphs is technically feasible while preserving the privacy of each participant's data.

Each component has been demonstrated. Their integration — particularly real-time generative inference over encrypted cognitive profiles at the latency learning demands — remains an engineering challenge, which is precisely why this requires Moonshot-level investment rather than incremental funding.

What prevented these pieces from closing the gap was not any single missing capability. It was the absence of integration. Generative AI is the catalyst.

Before large language models, a deep learner profile was inert data — rich but unable to generate a learning experience from what it knew about a person. Before edge AI reached its current performance, privacy-preserving computation could not reach learners without stable connectivity. Before federated learning matured, a knowledge graph could not be maintained across nations without centralizing the data in one jurisdiction. The pieces were puzzle fragments in separate boxes. Generative AI is not a new piece. It is the table on which the puzzle can finally be assembled.

But they must be assembled together. A deep learner profile without data sovereignty is a surveillance instrument. Data sovereignty without trustee governance is a vault without an access policy. A knowledge graph without a learner profile is a map without a position. Real-time generation without a validated knowledge graph produces fluent nonsense. Hub infrastructure without offline capability excludes the very population it claims to serve. Each sentence names a failure mode that is independently verifiable — and each explains why the intuitive response, *take the best existing platform and regulate privacy onto it*, cannot work.

The complete architecture — eight interdependent layers, nine design principles, the full dependency graph — is described in a companion paper (Pochmann 2026). The point here is structural: piecemeal assembly does not produce this system. It produces a more expensive version of what already fails.

The differentiation from the most capable actor sharpens at the level of category, not capability. LearnLM achieves real Socratic tutoring. Its measured learning gains are genuine. Its deployment at scale is an engineering accomplishment that deserves honest acknowledgment. What it does not generate: persistent learner profiles that accumulate across years. Simulation environments where the learner changes variables and watches forces respond — computed by deterministic physics engines, not rendered as text descriptions of what would happen. Process-based competency evidence that emerges continuously from the act of learning itself. These are not engineering limits — the structural constraints are defined above. The governance of the knowledge graph — particularly the boundary between universal scientific consensus and culturally sovereign content — is the most politically complex element of the architecture, and the feasibility study must address it directly. Google generates explanations *about* physics. This architecture generates environments *in which* one experiences physics. That is not a quality judgment. It is a category difference — the difference between reading about swimming and being in the water. What is at stake is who holds the blueprint of how a mind learns. Only trustee infrastructure can make the promise architectural.

The system complements teachers. Where teachers exist, it handles content delivery and formative assessment — returning to teachers what bureaucracy and class size have taken: the time to do what only humans can. Mentorship. Moral judgment. The moment a student needs not an explanation but a person. In the Global South, the calculus is different: 44 million additional teachers are needed by 2030 (UNESCO 2023), and the gap is widening. For children who have no teacher and foreseeably will not have one, the alternative to this system is not a teacher. It is nothing.

The strongest pedagogical objection remains: a system this responsive risks creating dependency. A learner who has never navigated ambiguity, endured a poorly structured resource, or decided alone what to study next has never built the self-regulation the world demands. The architecture answers through design, not disclaimer. With growing competence, the Mentor deliberately withdraws. It reduces the polymorphism that once translated every concept into the learner's native mode. It pulls back scaffolding. It introduces deliberate friction — problems presented without the tailored entry point, materials left unstructured, the profile temporarily suspended so the learner works with generic resources and self-directs. The principle is grounded in established learning science — Bjork's research on desirable difficulties, Kapur's work on productive failure — which demonstrates that strategic withdrawal of support builds deeper understanding than continuous scaffolding. Calibrating the withdrawal curve — when to reduce scaffolding, how fast, with what fallback if the learner is not ready — is an open research question that the feasibility study must address. The goal is not a learner who needs the Mentor forever. It is a learner who outgrows it. Early adaptation opens access. Later withdrawal builds autonomy. A mind that can only think in its own cognitive dialect has acquired knowledge but not resilience — and the world does not adapt to you.

Three scenes. A girl in Rwanda who learns what forces are by feeling them break her bridge. A man in Thessaloniki whose path to a new career runs through who he already is, guided by a profile no one else can see. A teenager in São Paulo who discovers that the system's deepest function is not to isolate her with a perfect feed but to connect her with the exact human whose thinking completes her own. Each scene is a window into what education becomes when the architecture serves the learner instead of the platform. The components that underlie each scene exist — physics engines, generative AI, edge hardware, federated learning. Their integration into a coherent system, particularly on low-cost hardware in low-connectivity environments, is part of what the Moonshot must achieve; if it were already solved, it would not need Moonshot funding. The scenes show the destination. The feasibility study determines the route. What connects them is the structural requirement: the learner, not the platform, holds the blueprint.

Why Europe — and Why Now

The MFF 2028–2034 proposal lists eleven Moonshot projects. Fusion energy. Quantum computing. Clean aviation. Next-generation AI. Data sovereignty. Automated transport. Regenerative therapies. Space economy. The system that trains the scientists for fusion, the engineers for quantum computing, the physicians for regenerative therapies — the floor on which every other Moonshot stands — is absent. Eleven doors. No floor.

The architecture requires, at its foundation, that no entity other than the learner can access their cognitive profile. This single requirement eliminates two of the three global powers capable of building at scale. The US technology sector operates under business models that treat learning data as a platform asset; the CLOUD Act allows US authorities to demand data from US companies regardless of where the servers stand. China operates under a state model that treats learning data as a governance instrument; the Belt and Road Initiative spans 146 to 150 countries, its Digital Silk Road component exports smart-city infrastructure to more than 100, Luban Workshops and Confucius Institutes build what researchers call epistemic communities, and the models that emerge carry hard-coded restrictions reflecting domestic political red lines.

Europe remains. Not by aspiration — by elimination. It is the only actor that combines a regulatory framework making data sovereignty a legal baseline, existing compute infrastructure, a seven-year budget mechanism capable of sustained commitment, and institutional precedent for building sovereign infrastructure through multilateral cooperation. A coalition of middle powers or a UNESCO-led initiative could contribute — the CERN model anticipates exactly this: European initiation leading to global participation. But initiation requires an actor with both institutional capacity and the political will to move first.

The obvious objection: Europe's track record on large-scale digital infrastructure is mixed. Gaia-X was supposed to deliver European cloud sovereignty and has largely stalled, captured by the hyperscalers who joined the consortium. The European Health Data Space is years behind schedule. But Gaia-X failed because it tried to compete with hyperscalers on their own terms — general-purpose cloud infrastructure. This project does not compete with Google on general-purpose AI. It builds specific-purpose infrastructure for a domain where the institutional requirements — data fiduciary, knowledge graph governance, offline operation, progressive de-adaptation — are structurally incompatible with the platform model. The CERN analogy is

precise: CERN did not try to out-fund Brookhaven. It built something Brookhaven's institutional form could not build.

Europe has done this before. Galileo was proposed not because GPS was technically inadequate, but because critical navigation infrastructure should not depend on a single foreign power's military system. Costs escalated to approximately €22 billion over twenty years. At multiple points the project was declared nearly dead. Each time, member states voted to continue. Civilian-controlled satellite navigation now serves nearly four billion users. The lesson is not that Europe builds smoothly. It is that when Europe decides critical infrastructure must not depend on a foreign power, the infrastructure gets built — even when the path is difficult.

The sovereignty concern is not abstract. More than 70 percent of European cloud infrastructure runs on US hyperscalers. The CLOUD Act allows US authorities to demand data from US companies regardless of where the servers stand. When the International Criminal Court's chief prosecutor was locked out of his Outlook account following US sanctions, the institution replaced Microsoft Office entirely. Now consider what a sovereign education architecture holds: not spreadsheets but how a person thinks, where they struggle, which memories matter. Storing this on CLOUD Act infrastructure means delegating cognitive sovereignty over a generation to a foreign jurisdiction.

The EU built that chain link by link — chip sovereignty, data sovereignty, digital sovereignty. Each was a response to dependence on external architectures. The chain has a missing link. Cognitive sovereignty is the logical next step in a sequence the Commission itself defined. Not completing it is not caution. It is incoherence.

The architecture is described. The precedent is established. What remains is the political arithmetic — and it is more concrete than the word "Moonshot" suggests.

The economics compress into two facts. For two hundred years, Baumol's cost disease meant every additional student required an additional teacher — costs grew linearly with enrollment. This architecture breaks that logic: marginal cost for an additional learner approaches zero. And the cost of the status quo is not stagnation but hemorrhage — the World Bank estimates $21 trillion in lost lifetime earnings for the current generation; Hanushek and Woessmann calculated that raising OECD PISA scores by 25 points over twenty years would yield $115 to $260 trillion in GDP gains. The EU invested €70 billion through the Recovery and Resilience Facility. PISA scores declined.

The political window is concrete and closing. The MFF 2028–2034 trilogue is underway. Final agreement is expected by end of 2027. Once agreed, budget

lines are fixed for seven years. Horizon Europe proposes €175 billion, the European Competitiveness Fund's Digital Leadership Window adds €54.8 billion, Erasmus+ stands at €40.8 billion. Education is absent from the Moonshot list. Every year of delay is also a year in which the commercial installed base grows by tens of millions, making the alternative harder to build and lock-in deeper to reverse.

Technological windows and political windows rarely align. When they do, the alignment is brief — the technology matures, the budget cycle closes, and the installed base calcifies. The generation entering school this year will learn inside whatever architecture exists when they arrive.

Three commissioners' portfolios converge on this project. Mînzatu holds Skills and Education — Erasmus+ at €40.8 billion in the proposed MFF, with 2026 calls allocating €54 million across ten policy experimentation topics, including AI ethics and STEM. Zaharieva holds Research and Innovation, with Horizon Europe at €175 billion. Virkkunen holds Tech Sovereignty — the AI Act, digital sovereignty, the infrastructure agenda. The institutional bridge exists at the personal level. What does not exist is the project that walks across it. The Commission has established inter-service mechanisms for projects spanning multiple DGs — the European Chips Act required coordination across DG CNECT, DG GROW, and DG RTD through an inter-service group and joint communication. Education infrastructure requires the same: a lead DG, a cross-DG mandate, and a budget line that reflects the project's infrastructure character rather than its educational label.

The demand is already being articulated — in fragments. UNESCO's Giannini warns that commercial AI risks eroding the human practices education depends on. UN Special Rapporteur Shaheed demands regulation and acknowledges the UN has no technical plan. More than fifty African states signed the Africa Declaration on Artificial Intelligence in Kigali, demanding sovereign computing infrastructure. Each sees a piece. No one has assembled them. The technical capacity for a first phase exists: Apertus, released by ETH Zürich, EPFL, and CSCS in September 2025, delivers 8-billion and 70-billion parameter open models trained across more than 1,000 languages, EU AI Act-compliant. Nineteen AI Factories are selected across Europe, five Gigafactories in planning. For a pilot, Europe has enough. For planetary scale, no single actor's infrastructure is sufficient — not Europe's, not America's, not China's. That is the point of a Moonshot.

The First Step

December 1949. Louis de Broglie stands before the European Cultural Conference in Lausanne and proposes what no single nation can afford: a laboratory for particle physics at the scale the science demands. The continent's best physicists have already left — Brookhaven is running, American budgets are growing, and Europe's remaining scientists see a central laboratory as a threat to their own funding. Governments hear "nuclear research" and think of the bomb. François de Rose, the French diplomat who carried the proposal through European capitals in 1949, later recalled: "We therefore failed in our mission. However, the idea was now on the table."

Six months later, Isidor Rabi introduced a resolution at the fifth UNESCO General Conference authorizing support for regional research laboratories. The resolution passed. In February 1952, eleven countries signed the provisional council agreement. In September 1954, CERN existed. Five years from idea to organization. Today: twenty-five member states, cumulative Member State contributions spanning seven decades, more than 12,000 visiting scientists from over eighty countries. In 1989, Tim Berners-Lee invented the World Wide Web at CERN — a side product of basic research infrastructure. CERN makes the fundamental structure of matter researchable. This project makes the fundamental structure of human learning researchable — and uses it simultaneously.

A continent whose best minds drift toward better-funded powers. National budgets insufficient for the scale the problem demands. Scientists worried about resource cannibalization. Governments uncertain whether the political risk of acting outweighs the political risk of not acting. An idea too large for any single country — and a small number of people who see that the idea, once on the table, cannot be responsibly ignored. The form followed the decision: a provisional council under UNESCO before CERN had legal existence. Three scientists — Rabi, Auger, Amaldi — initiated it. Governments joined because the alternative was to watch the capability leave.

The cost question inverts under examination. The World Bank estimates $21 trillion in lost lifetime earnings for the current generation; Hanushek and Woessmann calculated $115 to $260 trillion in GDP gains from a modest 25-point PISA improvement. No single project eliminates that loss. But no project addresses it at scale without the infrastructure layer this Moonshot would build. These figures measure the magnitude of what is at stake — the economic cost of the problem this infrastructure addresses — not the return on a single investment. The sharper inversion is structural: the MFF proposal funds eleven Moonshot projects. The system that produces the people for all eleven is not among them.

The demand is specific. Education as a twelfth Moonshot project in the MFF 2028–2034 — not as an extension of Erasmus+, but as infrastructure on the level of fusion energy and quantum computing. The budget category exists. The architecture is described in a companion paper (Pochmann 2026). The first step is not to build the system. It is to commission the feasibility study — the same step that moved CERN from an idea at a cultural conference to a provisional council in under three years. A feasibility study is a decision to look. Refusing to look is also a decision.

In the early 1980s, each European country operated its own mobile telephone standard. Calls stopped at borders. The United States chose competing proprietary standards; the market fragmented for a decade. Europe chose differently: GSM, an open standard any manufacturer could implement, any country could adopt. The decision was political, not technological. GSM became the global standard. A continent that could not agree on electrical plugs agreed on how radio waves would carry voices — and that agreement shaped global telecommunications for a generation.

AI-powered education is arriving with the same inevitability. The cost curves, the model capabilities, the connectivity trajectories all point in one direction: every learner on Earth will have an AI tutor within a decade. What remains undecided is the architecture — and the cognitive lock-in that began with a feature update in 2025 deepens with every year the question stays open. Whether that tutor operates as sovereign infrastructure or as a platform product is not a technology question. It is a question of who holds the record of how a generation learned to think.

> **One path produces citizens who own their learning. The other produces users who rent it.**

References

Africa Declaration on Artificial Intelligence (2025). Signed at the Global AI Summit for Africa, Kigali, Rwanda, April 4, 2025.

Associated Press (2025). “Microsoft blocks ICC prosecutor’s email following US sanctions.” May 2025.

Baumol, W.J. & Bowen, W.G. (1966). Performing Arts: The Economic Dilemma. Twentieth Century Fund, New York.

Bjork, R.A. & Bjork, E.L. (2011). “Making Things Hard on Yourself, But in a Good Way: Creating Desirable Difficulties to Enhance Learning.” In M.A. Gernsbacher et al. (Eds.), Psychology and the Real World, pp. 56–64. Worth Publishers.

ETH Zurich, EPFL & CSCS (2025). “Apertus: a fully open, transparent, multilingual language model.” Press release, September 2, 2025.

European Commission (2025). A Dynamic EU Budget for the Priorities of the Future – The Multiannual Financial Framework 2028–2034. COM(2025) 390 final. July 16, 2025.

European Commission (2025). Commission Work Programme 2026: Europe’s Independence Moment. COM(2025) 870 final. EUR-Lex: 52025DC0870.

Google (2025). LearnLM: AI for Learning. Official product page. Google for Education.

Google DeepMind (2025). LearnLM Tutoring Study: UK Controlled Trial Results. arXiv:2512.23633.

Giannini, S. (2023). “Harnessing the Era of Artificial Intelligence in Higher Education: A Primer for Higher Education Stakeholders.” UNESCO, Paris.

Hanushek, E.A. & Woessmann, L. (2010). The High Cost of Low Educational Performance: The Long-Run Economic Impact of Improving PISA Outcomes. OECD Publishing, Paris.

Kapur, M. (2008). “Productive Failure.” Cognition and Instruction, 26(3), 379–424.

OECD (2023). PISA 2022 Results (Volume I): The State of Learning and Equity in Education. Paris.

OECD (2026). OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education. OECD Publishing, Paris. https://doi.org/10.1787/062a7394-en

Open Europe (2010). Galileo Programme Cost Estimate. London.

Pochmann, M. (2026). A Generative Education Architecture for Planetary-Scale Personalized Learning. Companion Paper. https://doi.org/10.5281/zenodo.18759134

Shaheed, F. (2024). The Impact of the Digitalization of Education on the Right to Education. Report A/79/520, United Nations General Assembly, 79th Session.

Stanford University Human-Centered AI Institute (2025). AI Index Report 2025. Stanford, CA.

Synergy Research Group (2025). European Cloud Infrastructure Services Market Analysis. Reno, NV.

UNESCO (2022). Transforming Education Summit Discussion Paper.

UNESCO (2023). Global Report on Teachers: Addressing Teacher Shortages. Paris.

UNESCO Institute for Statistics & GEM Report (2025). 2025 SDG 4 Scorecard: Focus on the Out-of-School Rate.

WinBuzzer (2025). “International Criminal Court Drops Microsoft 365 for European Open-Source Suite.” October 31, 2025.

World Bank, UNESCO, UNICEF (2022). The State of Global Learning Poverty: 2022 Update. Washington, DC.